

USING HIGH PERFORMANCE COMPUTING TO REALIZE A SYSTEM-LEVEL RBDO FOR MILITARY GROUND VEHICLES

David A. Lamb
David Gorsich
U.S. Army RDECOM-TARDEC
Warren, MI 48397
(Kerberos Principal: HPCMP)

Introduction

Reliability Based Design Optimization (RBDO) is a process of optimizing under uncertainty to obtain a reliable (in the probabilistic sense) optimum for a design, which is robust under the expected variability inherent in realizing the design. In this case, we are optimizing the design of a ground vehicle to reduce weight while maintaining or improving durability. Like any optimization, it is best done on a system level. When optimizing under uncertainty, considering a large number of sources of variability makes the optimization method more robust.

Objective

The objective is to improve the RBDO process, expanding from component optimization to system optimization of a ground vehicle, consider more sources of uncertainty and use a multi-physics and multi-scale approach. This requires greater computing power than has previously been applied to such optimizations.

Methodology

The massively parallel computing power of the Department of Defense (DoD) High Performance Computing (HPC) systems is used to simultaneously optimize multiple components which interact with each other in a mechanical system. Specifically, we have a subsystem of a military ground vehicle, consisting of at least four components and we are simultaneously optimizing all the components of that subsystem using RBDO methods. We do not simply optimize one component at a time, sequentially, and iterate until convergence. Instead, we simultaneously optimize all components together. This can be done efficiently using a parallel computing environment.

Results

The speed up realized by parallelizing enables the jump from component level to system level optimization, the addition of multiple physics-of-failure in the analysis, and consideration of vastly more sources of uncertainty than could be achieved from serial computing.

Significance to DOD

In order to reduce the weight of ground vehicles while maintaining or improving durability, RBDO is a key enabler for DOD. The only efficient way to accomplish this goal is to parallelize the method.

1. Reliability Based Design Optimization (RBDO)

In order to reduce the weight, improve the reliability and increase the lifetime of a complex mechanical system, such as a ground vehicle, we need to know how to use modeling and simulation (M&S) to assess durability and maintenance. Then we can optimize the system with weight as the objective to minimize, but with the durability as a constraint. This should allow us to both reduce the weight and increase the availability of the ground vehicle. This is what Reliability-Based Design Optimization (RBDO) attempts to do.

The modeling of complex mechanical systems for durability is a challenge, however, for several reasons. The durability of a complex mechanical system is affected by many different ways failure can occur, including (but not limited to) overload, fatigue, thermal, corrosion, and wear. Significantly, these factors are not really independent of each other, and there can be significant coupling between thermal loads, corrosion and fatigue-inducing mechanical loads for failure of systems. Therefore, to get the most accurate assessment, we can't treat these different "physics of failure" separately, but must combine them into a single model/simulation.

Another major complication with this is that

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 14 JUL 2008		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Using High Performance Computing to Realize a System-level RBDO for Military Ground Vehicles				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) David A. Lamb; David Gorsich				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) US Army RDECOM-TARDEC 6501 E 11 Mile Rd Warren, MI 48397-5000				8. PERFORMING ORGANIZATION REPORT NUMBER 18944RC	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S) TACOM/TARDEC	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) 18944RC	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES Presented at the Department of Defense (DoD) Hight Performance Computing (HPC) Users Group Conference, July 14-17, 2008, Seattle, Washington, The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 7	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

fatigue (to name a common physics-of-failure) is not a deterministic process, but rather is a stochastic one. Other stochastic elements are also important to the durability of the system, such as material impurities, geometric tolerances in manufacture, and usage history. It is very inappropriate to think that we can ever get a deterministic result from the simulation for anything in the area of fatigue lifetime or even for general durability. Stochastic methods are extremely important to capture all the sources of variability in the system: material, manufacturing, operating environment, maintenance history, and differences in operator behavior. The answers will all be phrased in terms of probability distributions.

Then, as a final complication, the optimization we wish to perform is not going to an optimization based on deterministic constraints and objective functions, but rather is “optimizing under uncertainty” where the objective and the constraints are stochastic measures, and we need to define some level (an α -cut) of staying inside the constraints. For example, do we want a 90% chance that we don’t violate constraints, taking into account all stochastic variability, or would 95% or 99% be a better level of confidence? The optimum computed will be affected by the α -cut chosen.

1.1 The Scope of the Problem

Prof. K.K. Choi, of the University of Iowa, previously performed [Choi, 2001] a reliability-based optimization of the design for an A-arm on a current military ground vehicle, using no sources of uncertainty and only one physics-of-failure (fatigue). This was done using serial computing. He reported using 768 Finite Element Analysis (FEA) runs of small-sized models (30K – 200K DOF) and taking 3.55 days of compute cycles. This was just for a single component and a single physics.

While he demonstrated part of the method, he was criticized for optimizing a single component in isolation from the rest of the system. System optimization is best done by considering the system as a whole, and not simply iterating through an optimization one component at a time, hoping for convergence to a system optimum at the end.

Prof. Choi estimated that to do a full vehicle, his method would take at least 100 times the computing cycles, or 76,800 FEA runs and 355 days in serial mode. But, he reports, the FEA are all largely independent and could be done in

parallel. Utilizing 1,000 processors each capable of doing a single FEA run on a small-size model in serial, he projects that the turn-around time drops to below half a day.

1.2 Our Goal

We are planning for something even more ambitious, since we want to optimize a whole system (many components interacting with each other) while using four or five physics-of-failure and many sources of uncertainty requiring Monte-Carlo techniques. Estimates climb into the tens of millions of FEA runs of small-sized models, and hundreds of years of clock time if done in serial. Fortunately, there is no need to do this in serial, since most of the FE analyses are independent, and we can parallelize. Utilizing 10,000 processors to parallelize the FEA runs will keep the turn-around time below two weeks. To be useful in influencing the acquisition process, turn-around times longer than week are not acceptable. Unfortunately, we cannot immediately jump to using 10,000 processors, but will have start out more modestly and grow to that level.

2. THE METHOD

Two key features of this method are that it is physics-based, starting from first principles, rather than heuristic, and that it seeks to handle interactions between different components of the ground vehicle and different physics-of-failure on a non-heuristic basis. We are seeking methods to compute fatigue, thermal stress, corrosion and other causes of failure using physics-based equations as can be found in textbooks or handbooks, and do not want to depend on heuristically generated response surfaces or some other ‘rule of thumb’ based on statistical manipulation rather than physics first principles. We want to predict the reliability of the ground vehicle starting at the material level, working up through components, assemblies and subsystems to the system level, while having a good scientific basis for each step, rather than just a statistical basis.

Understandably, this takes a large amount of computing to accomplish. We parallelize at several different levels, including assigning different components to run on separate sets of processors and by configuring different physics of failure onto their own processors. This, however, is not perfect, as the problem requires

some coupling between the different physics and also between components. This coupling must be accounted for somewhere in the simulation. Still, with a scheme of dividing the problem up by parts of the vehicle, failure modes, and dealing the stochastic uncertainty using multiple processors, we plan to rely on the High Performance Computers (HPC) to accomplish the computational analysis.

The intended end-use of this method is to quickly and accurately generate a prediction of the reliability for a proposed design, so that this prediction can be used for trade-off studies or for optimization of the design. As such, the method must only use input which would be generally available during the design cycle when trade-off studies are performed. To actually have any influence on the final design, the prediction must be accomplished in a short amount of time, so the results are available for the next design iteration. We expect that unless a prediction can be made in a week, we will miss the opportunity to guide the design loop process toward greater reliability.

2.1 Massive Number of FEA Runs

The main idea that we are using is that the reliability analysis incorporates a large number of FEA analyses, most of which are independent. The greatest speedup in time to final answer will come from spreading the FEA runs across a large number of processors to be executed in parallel. This will require methods to break the large scale systems into lower scale ones, and methods to break apart different physics-of-failure into separate analyses loosely coupled with each other. An automated process for generating the necessary multiplicity for the Monte-Carlo technique to address the uncertainties will be needed. Finally, a method to consolidate the results back up to the system level will be required.

2.2 Course Grain versus Fine Grain Parallelization

We did a preliminary study to decide if there is an advantage to parallelize a single FEA run, or simply run a number of FE analyses in a serial fashion simultaneously. The results of this study showed that our typical FEA runs are not particularly large, but we need a lot of them run. Culling from the analysis of the A-arm done in [Choi, 2001], we estimate that a ground vehicle consisting of 100 components, using four physics-of-failure, and 100 Monte-Carlo points

for computing the stochastic distribution will require 30,720,000 finite element analyses (FEA) each in the range of 30~200k DOF. See figure 1 for an example of the process flow used to derive these numbers.

Thus, significantly more speed up could be achieved by carrying out a number of FE analyses simultaneously, rather than trying to make each FE analysis faster. Parallelizing by putting one FEA on each processor but running 1000 at a time counts more than spreading a 200k DOF FEA across 100 processors.

As it turns out, while this is a very good way to parallelize the method, it leads to a significant challenge for the project, as we will discuss later in this paper.

2.3 The Challenges

We expected to find several challenges in the computational process caused by the need to generate, coordinate, and finally consolidate the individual results. At the lowest level, we rely on native queueing software to coordinate scheduling many FEA runs onto the processors.

We did find a number of challenges. We needed work flow software so we scripted our own work flow control. This provided a challenge, but we overcame it and now have in-house scripted work flow control for all further in this research.

We also encountered a challenge obtaining the base data needed for the study, particularly in the area of uncertainty distributions for the material properties of the steel in the part being studied. This is discussed further below.

3. THE PROJECT

We made the first set of runs in the May-June 2007 timeframe on the HPC systems located at U.S. Army RDECOM-TARDEC in Warren, MI. We describe here the results seen in these runs.

We analyzed the lower driver's side A-arm from another military ground vehicle. (See figure 2 for the part analyzed.) We set up an optimization to reduce the weight of the design and to improve fatigue life. We chose this part because it was very similar to another study done using serial processing earlier, and there was enough data available for this vehicle and this part to serve as a test case.

We wanted to do a multi-scale, multi-physics analysis of a subsystem, but being

limited on resources we could bring to the pilot project, we found that the only way to get anything run was to be more modest in our immediate goals. The pilot project was restricted to only did a single component and a single physics-of-failure.

3.1 The Computer Hardware

Three computer systems were used for this project. The first was eight 1.3 GHz processors, 8 Gbytes memory and 72 Gbytes local disk space. The second was equipped with 24 MIPS processors, 24 bytes memory and 72 Gbytes local disk space. The third was implemented with 32 MIPS processors, 32 Gbytes memory and 36 Gbytes local disk space. All three are part of the TARDEC HPC System. TARDEC HPC has ties with the DoD HPC Modernization Program.

3.2 Reliability/Fatigue Analysis software

We used fatigue analysis software, design sensitivity software and reliability-based design optimization code. All three were ported to run in the TARDEC HPC environment.

In addition to these, we made use of numerical analysis software. This was used primarily to perform the optimization in the loop.

3.3 Finite Element Analysis solver

We needed extensive use of a finite element analysis solver. To accomplish significant parallelization of the method, we required multiple copies of an FEA solver be running on different processors, solving variations of the same analysis, in parallel. We found that most vendors of FEA code treat this situation as requiring a license for each solver run. So, to run on sixteen processors required having sixteen licenses, and to run on one hundred processors would have required one hundred licenses.

3.4 Parallelization and work flow control

RBDO demands multiple reliability analyses for a given design. In the pilot study, we refined the method to allow that reliability analyses be performed only for the active/violated probabilistic constraints. These were executed in a parallel manner on the HPC system. By eliminating the non-active constraints, we reduce the computation burden. Thus, by this reduction, fewer processors are needed to parallelize the

entire process of reliability analysis. The parallelization has been successfully tested using LSF on the TARDEC HPC.

3.5 Preprocessing software

We required multibody dynamic analysis of the whole vehicle to obtain loads for the fatigue analysis. This dynamic analysis was performed in a preprocessor step. This was not done during the parallelization stage, and the same loads were used throughout the entire pilot run. The dynamics software was just for preprocessing the dynamics loads.

We also used meshing software for creating the original mesh on the part we were analyzing. This was done once in a preprocessor step. The FEA software was run in a preprocessor step to determine 'hot spots' and pre-configure the fatigue solving step.

3.6 Results of scalability study

The RBDO parallelization was tested out by conducting a scalability study using different combinations of processors from 1 to 32, licenses of the FEA solver software from 1 to 16, and other settings. The throughput times were then compared to get a sense of how the problem throughput would be affected by larger numbers of processors, FEA licenses, etc. The results of this study are provided here.

For the scalability study, 22 experiments (20 training runs and 2 test runs) were designed with different numbers of FEA solver licenses, processors, and constraints. We noticed that a dependence of the parallel runtime (PR) on the number of FEA licenses occurs when the number of licenses is less than the number of processors and individual constraint runs are forced to wait for a license to become available. For the MPP search, finite element analysis accounts for about 22% of computational time in a serial run. So the number of licenses has a large effect on the parallelization of the process, but does not completely control the degree of parallelization. For the 20 training runs, 1, 2, 4, 8, and 16 licenses, 1, 8, 15, and 30 processors, and 15 and 30 constraints were used. Not all possible combinations made sense for a run. In particular the number of processors should be greater or equal to the number of licenses, else there are unused licenses. We had a slight violation of this rule for two of the runs, since configuring those runs for all the available 16 licenses was more

natural. Also the number of constraints should be greater or equal to the number of licenses and the number of processors; else there are unused licenses or processors. Again a slight violation of this rule is present in two of the runs. Finally two test runs were performed as a “sanity check” on using the training runs in a predictive way. Please see Figures 3 and 4 for two interpolation surfaces of runtimes based on numbers of processors and FEA licenses used.

4. FOLLOW-ON STUDY

The results of the pilot experiment were promising enough to have generated a follow-on study to extend RBDO into more areas. We want to do a multi-component optimization at a system-level, as well as increasing the number of sources of uncertainty. We want to extend, also, into physics-of-failure beyond fatigue, such as thermal and corrosion.

A study is underway now, in the model development phase, for a six body RBDO of a frame and powerpack for another military ground vehicle. Of the six bodies in the problem, only four are “active” in the sense that we are allowed to modify them in the optimization. The other two are “constraint” bodies which we must have in the system, but leave unchanged. We are greatly increasing the number of sources of uncertainty, and we are adding a thermal component to the reliability, in addition to fatigue.

This study is expected to be executed on the TARDEC HPC system in the Fourth Quarter of FY08. The new TARDEC HPC is expected to be available in August. At the current rate of model development, we fully expect to be executing by September 2008.

5. THE PAYOFF

When talking about reliability, it is important to consider ‘total lifecycle cost’ as the relevant measure. This is because designing to increase reliability often costs extra at the front end during research, development, design and manufacturing phase, but realizes savings during the Operations and Sustainment phase of the life cycle due to reduced costs to keep the vehicle available. To understand the value added by the increased reliability, the key is to balance the added up front costs against the deferred savings. In other words, we need to look at total cost across the entire life cycle of the vehicle.

Following the law of diminishing returns, the projected savings from improved reliability is often based on the starting level of reliability in the system. If a fleet is showing low reliability before efforts begin, then a large cost savings due to improved reliability is possible, but it is hard to realize great savings when starting from a fleet of very reliable vehicles. Based on current data, it appears that improved reliability in Army ground vehicles has a potential for respectable cost savings.

Total savings will also be a function of the number of similar vehicles in the fleet based on the improved design. It is obviously easier to realize large cost savings from improving the reliability of a design with 10,000 fielded vehicles than improving a design that only fields 50 vehicles. Still, once methods are developed to improve the reliability of a design, and the cost to develop the methods is recouped from improving the design of a few vehicles, the same methods will still be available to use on all other vehicle designs with minimal added cost. The key, therefore, is to apply the new methods to a few systems where the development costs of the new methods can be quickly recouped, and then deliver to the Army a ‘paid for’ tool to improve the reliability for other Army platforms.

There is potential for tens of millions of dollars in total life cycle cost savings for a fleet of a single ground vehicle design due to improved reliability designed from the beginning. These savings will be spread across the whole life cycle as well as across the fleet of similar vehicles. If this method can be used to improve the design of just ten future vehicles, with various sizes of fleets and various results of reliability improvement for each, the method could potentially lead to savings that provide an excellent return on investment in the RBDO development. Even just one vehicle design will more than repay the costs of developing and implementing the method, based on modest reliability improvements to the design from the use of this tool.

CONCLUSIONS

While the Army struggles with the reliability of its current and future fleets of ground vehicles, there is a technology gap which can be bridged with RBDO. We want to make it a good tool, one based on physics and not heuristics, and one that considers system level reliability with interactions between components and between failure modes captured. This requires the

massively parallel environment of HPC to be realized quickly enough to impact the design loop. We are working to build this technique, make it multi-physics, multi-scale and non-heuristic. As this project progresses, we will add additional complexity to the models and generate predictions that encompass the true range of reliability.

REFERENCES

- Davis, TP., 2004: "Science engineering and statistics"; Henry Ford technical Fellow for Quality Engineering; Ford Motor Co.; <http://www.timdavis.co.uk/research.htm>.
- Kokkolaras, M., Mourelatos, Z. P., and Papalambros, P. Y., 2004: "Design Optimization of Hierarchically Decomposed Multilevel System under Uncertainty", Proceedings of the ASME 2004 Design Engineering Technical Conferences, Salt Lake City, Utah, September 28 - October 2, 2004, DETC2004/DAC-57357.
- Youn, B.D. and Choi, K.K., 2004: "Selecting Probabilistic Approaches for Reliability-Based Design Optimization," AIAA Journal, Vol. 42, No. 1, 2004, pp. 124-131.
- Choi, K.K. and Youn, B.D., 2001: "Hybrid Analysis Method for Reliability-Based Design Optimization," 27th ASME Design Automation Conference, September 9-12, 2001, Pittsburgh, PA.
- Socie, Darrell F., 2005: Probabilistic Aspects of Fatigue, <http://www.fatiguecalculator.com/pdfs/ProbabilisticFatigue.pdf#search=%22SAE%20950X%20deviation%22>

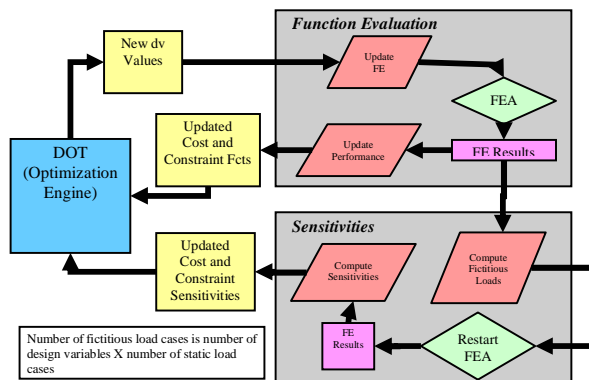


Figure 1. Example of method described.

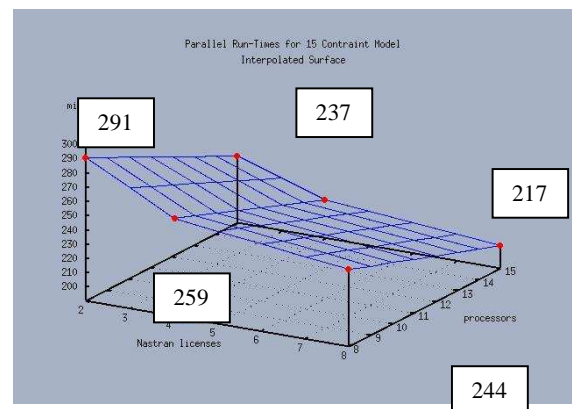


Figure 4. Scalability Interpolation Surface for 15 constraints.

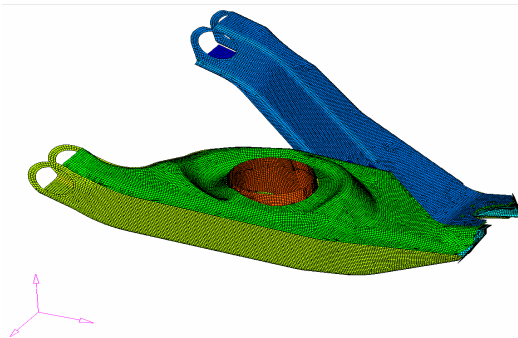


Figure 2. Lower A-arm.

Figure 3. Software loop diagram.

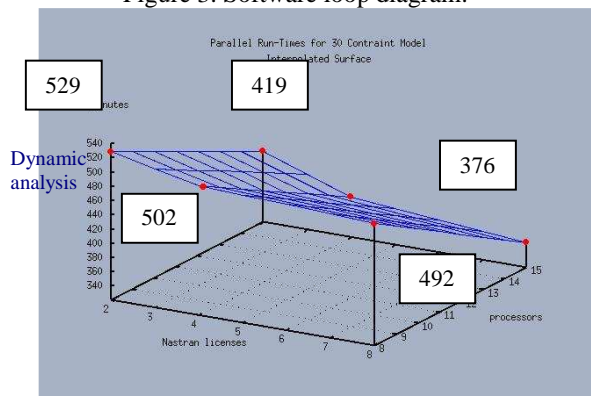


Figure 3. Scalability Interpolation Surface for 30 constraints.

354